AEIC GRID ADVANCEMENT PROGRAM · CHALLENGE BRIEF

Flexible Load Forecasting for a Nonlinear Future

Challenge 07. The operating problem, why it is hard, and what the cohort will work.

Working session	Tuesday, February 9, 2027, 2:00 p.m. ET
Cadence	Second Tuesday of every month, 2:00 p.m. ET
Eligibility	AEIC members and any electric utility
Join	https://aeicopscenter.org/grid-advancement
Program	AEIC Center for Operational Excellence

The problem

The utility planning model runs on peak. One hour, the highest hour, and everything downstream gets sized against it. Capacity, capital, interconnection, resource adequacy, the rate case. That worked for decades because the load curve was boring. It moved a little, it moved slowly, and last year told you most of what you needed to know about next year.

That is over. The demand curve went nonlinear, and peak-load thinking breaks when it does.

Rooftop solar moved the net peak into the evening and hollowed out the middle of the day. Storage moves load on purpose, on an economic signal, which means the load shape is now partly a decision somebody else is making. Electric vehicles pile onto the hours residential load already occupies. Heat pumps are flipping summer-peaking territory toward winter constraints. And a single data center request can land a decade of organic growth on one substation in eighteen months.

None of that is visible in the single highest hour of the year. It is visible in the other 8,759.

Why this challenge

Because a forecast is not a report. It is the first domino.

Every capital dollar, every interconnection study, every non-wires alternative, every rate case, every hosting-capacity number traces back to a load forecast somebody built. When the forecast is wrong, the error does not stay in the forecast. It shows up years later as a substation that never got used, or a feeder that got overwhelmed, or a flexible resource that would have been cheaper than steel and never got evaluated because nobody could see the hour where it mattered.

And the errors are asymmetric now. Overbuild, and you are defending a stranded asset in front of a regulator who wants to know why customers paid for it. Underbuild, and you are defending a reliability event in front of the same regulator. There is no safe direction to be conservative anymore, which is a genuinely new condition for this discipline.

Where utilities actually are

Almost everyone is somewhere on the same three-step path, and the public version of where any of us are is tidier than the working version. That is not spin. That is what happens when you have to present a work in progress to people who need a number.

Peak-centric planning, dressed up. Annual or seasonal peak, escalated by an economic growth rate, with a DER adjustment bolted onto the end. It is the workhorse. It is also a top-down model, and top-down models cannot see a feeder.

AMI data that exists but is not in the forecast. Nearly every utility now has interval data for most of its customers, which was itself an enormous lift. Few have turned it into validated 8760 customer-class curves a planner will actually sign. The data is sitting there, in a system the forecasting team does not own, and the cleanup it needs has never been anyone's funded assignment.

A large-load queue nobody trusts. Every planning shop now has a spreadsheet of speculative data-center requests, most of them shopping the same load across multiple utilities. The forecast has to say something about that load, and there is no honest analytic method for a number that is somebody else's option to exercise.

None of this is wrong. It is what happens when the tools were built for a world where the past predicted the future, and that world ended faster than any planning cycle could turn.

The hard part

The math mostly exists. Hourly time-series forecasting, bottom-up DER-aware modeling, probabilistic scenario generation, weather-normalized customer-class disaggregation. All real, all published, all running somewhere. The hard part is these four things.

1. History stopped being a training set. Every conventional forecasting method is a way of extrapolating from what already happened. Solar adoption, storage dispatch, and EV uptake do not extrapolate. They move on policy, price, and consumer behavior, and they break trend without warning. Data-center load does not arrive as a trend at all. It arrives as a step. You cannot regress your way to a step change, and that is where most forecast error is now being manufactured.

2. A forecast is now a scenario set, and nothing downstream will take one. The honest output of modern forecasting is a probability distribution across a range of hourly futures. The capital plan wants a number. The interconnection study wants a number. The rate case very much wants a number. So somebody, usually quietly, collapses the scenario set back into a single line, and every bit of information the new method bought gets thrown away at the handoff. That is a governance problem, not a modeling problem, and it falls in the space between the forecasting team and the capital committee, where no one is assigned to stand.

3. The precision is bought with unglamorous data quality. An 8760 forecast at feeder resolution requires that you know which meter hangs off which transformer, which customer is really in which rate class, and which interconnected DER is actually energized. Those three facts live in three systems, they disagree with each other, and reconciling them is rarely funded and rarely owned. You cannot build a defensible customer-class curve on a customer class that is wrong.

4. The intelligence sits with planners and the decisions get made everywhere else. The forecasting team is small, technical, and three layers from the capital committee. Operators need to know what tomorrow's net load looks like. Executives need to know whether to commit to the interconnection. Regulators need a method they can approve. A scenario set that only the forecasting team can interpret is not yet a planning capability, and getting it there is an organizational job, not an analytic one.

If you need one sentence to carry into a capital committee: **there is no safe direction to be conservative anymore. Overbuild and we defend a stranded asset; underbuild and we defend an outage. A single-number forecast hides that trade from you, and you should not let it.**

None of those four barriers is local to you. In some form, all four show up in all ten challenges in this program: the model that cannot see what it is being asked about, the data reconciliation that is rarely funded, the work that falls between two functions, and the handoff where the honest answer gets flattened into the confident one. If your list looks like this one, you are not behind. You are in the same room as everyone else.

And "we do not know what that load is going to do" is a legitimate answer here. It is a finding. A range you can defend is worth more than a point you cannot.

What the cohort will work

Not "which model is best." The room works what actually blocks adoption:

- How are you forecasting a large load that has not signed anything, and what did you tell your capital committee about it?
- What did it take to get validated 8760 customer-class curves out of AMI data, and who owned the data cleanup?
- Has anyone gotten a probabilistic forecast, presented as a range, through a regulator? What survived?

- Where did bottom-up feeder-level forecasting genuinely beat the top-down number, and where did it not justify the effort?
- What did you try that did not work? That is the most valuable hour in this program, and it is the one you can only get from peers who have actually tried it.

What to bring

You do not need a solved problem. You need an honest one. Come with:

- One forecast you got wrong, and what it cost
- What your current forecast is actually built from, and at what resolution
- How you are handling speculative large-load requests today
- One thing you tried that did not work

If the honest answer to any of those is "we do not know," bring that. It is a finding, not a failure, and you will not be the only one saying it.

What you get

- This brief, and a short video briefing, before the session
- The Flight Plan, a working pilot scaffold, inside InnovationWorks
- The session itself, with the utilities living this problem
- A retrospective afterward: what the cohort tried and what happened

Join

Tuesday, February 9, 2027, 2:00 p.m. ET. Recruitment runs per challenge, so your planners and forecasters can come to this one without committing to all ten.

Open to AEIC members and to any electric utility. If you run a utility, you are welcome here.

https://aeicopscenter.org/grid-advancement
Knowledge of one is the knowledge of all.
Sources: this challenge was developed from research of public utility filings. The technology landscape for this challenge is maintained separately inside InnovationWorks.